Project name:
Reinforcement learning for graph theory

Project acronym:
/

Project duration:
July 1, 2024 – July 1, 2026

Funding program:
Bilateral projects for mobility of researchers between Serbia and PR China (2024-2026) funded by the Ministry of Science of the Republic of Serbia

Partner institutions:
- Serbia: Mathematical Institute SANU
- China: Nankai University, Center for Combinatorics

Serbian principal investigator:
Dragan Stevanović, Mathematical Institute SANU

Serbian research group:
Ivan Damnjanović, University of Niš

Chinese principal investigator:
Jiangdong Ai, Nankai University

Chinese research group:
Yongtang Shi, Nankai University
Hui Lei, Nankai University
Yandong Bai, Northwestern Polytechnic University, Xi'an
Susu Wang, Nankai University

Project description:
Reinforcement learning (RL) is a part of artificial intelligence methodology that has become very prominent recently due to its success in achieving superhuman levels of performance in playing go, chess, or good old video games. In RL an agent issues a list of actions to the environment, which then responds with a reward for each action. The agent on its side uses a deep neural network to learn how to maximize the reward obtained from the environment, so that RL is actually a particular sort of an optimization algorithm. While publicly RL has mostly been used in learning how to play games, Adam Zsolt Wagner [arXiv:2104.14516] has recently shown how it can be also used to construct graphs, objects from discrete mathematics, that have predetermined properties. In this case, the actions that the environment accept correspond to the edges of the graph (listed in some fixed order), and each 0/1 action determines whether the edge exists or not. The reward is withheld until the graph is fully constructed, and the final reward informs the agent whether the graph has the necessary properties (and in what amount). In this way one can construct counterexamples for open conjectures in graph theory or construct examples of graphs that could help obtain better insights when working on a new theoretical problem. Our goal here is to improve upon Wagner's approach in a number of ways: by implementing environments that construct more specific graph classes such as trees, signed graphs, uniform hypergraphs, etc.; by using more advanced RL methods such as proximal policy optimization or deep Q-networks instead of the old-fashioned cross-entropy method that Wagner used; or by enabling one to compute graph invariants in a much faster way. During this work we will also use the implemented framework to check for the existence of counterexamples for a number of open conjectures in graph theory.